New AI Rules for Colorado Real Estate

Imagine you're a property manager, right? You've got a stack of like 50 applications for a single rental unit. Oh yeah, a total nightmare to sort through manually. Exactly.

Yeah. So you run them through your shiny new highly expensive tenant screening software and it just instantly flags one applicant as high risk. So you know you deny them and move on.

Makes sense, that's why you bought the software. Right, but then a month later the Colorado Attorney General is knocking on your door basically threatening a deceptive trade practice lawsuit. And you're wondering why? Well because your software hallucinated a fake prior eviction.

The applicant found out and you completely failed to provide this highly specific legally mandated paper trail explaining the algorithm's decision. Yeah and that scenario, I mean that is the new reality in Colorado. It really is.

Welcome to our discussion today. Yeah. If you're operating in the Colorado real estate market and you want to you know stay ahead of the regulatory curve without drowning in legalese, you are absolutely in the right place.

For sure, because this is going to catch a shocking number of real estate professionals completely off guard. We are talking about a tectonic shift in how you're legally permitted to use technology in your business. Yeah, we're pulling from a massive stack of legal analyses for the topic today.

We've got sources from Alston and Byrd, Ballard, Spahr, Clark Hill, plus the actual legislative summaries from the Colorado General Assembly. And honestly we need every single one of those sources because the state legislature just completely pulled the rug out from under the tech industry. Yeah, they really did.

So the mission for our discussion today is to dissect Colorado's brand new AI legislation. That's Senate Bill 26189. We want to isolate exactly what it means for the real estate sector.

Like how does this actually impact real estate brokers and property managers? Right, the boots on the ground. Exactly. And here's the massive headline that you, the listener, need to internalize immediately.

This new law completely repeals and replaces the previous 2024 Colorado AI Act. Completely wipes it out. Yeah, the old rules are gone, dead.

There should be absolutely zero confusion for you moving forward. We are only talking about the new regime under 26189. And you know to understand the sheer scale of the administrative burden coming your way you really first have to understand why the legislature felt compelled to do this.

Taking the unprecedented step of entirely scrapping a major piece of legislation just two years after passing it. Right, because the 2024 law, SB 24205, that was technically the first comprehensive state AI law in the whole nation. It was, but operationally.

I mean it was a total disaster. I remember the uproar. But what was the actual like mechanical failure of that 2024 law? Why did the business community and even the White House call it burdensome? Well the fatal flaw was really its core philosophy.

It imposed this incredibly broad, really vague duty on businesses to use quote reasonable care to avoid algorithmic discrimination. Okay, which sounds good on paper I guess. Right, it sounds noble.

But in practice it required companies to conduct these massive nebulous impact assessments before they could even use an AI tool. Yeah, so if you were a mid-sized property management firm, how do you even begin to prove to the government that your third-party software doesn't have some underlying bias? I mean the technical demands were just impossible for a standard business to meet. So it was essentially demanding that like a real estate broker also act as a senior data scientist just to screen a tenant.

Precisely, and the legal backlash as you can imagine was immediate and intense. Yeah. An AI company called XAI actually filed a constitutional lawsuit in federal court to block the law claiming it violated the first amendment.

Wait, I need to stop you there. Okay. How does a tech company claim an AI regulation violates the first amendment? Like is a computer program exercising free speech? I know it sounds totally wild but it is a major legal argument right now.

The premise is that algorithms are essentially complex opinions expressed through code and the data they output is a form of protected speech. So forcing a company to alter how its algorithm speaks or restricting what it can output is, they argue, government compelled speech. That is fascinating.

Yeah, and to add fuel to the fire, the U.S. Department of Justice jumped in and joined the lawsuit against the state raising separate challenges under the equal protection clause. Wow, and just to be clear for everyone listening, we are purely reporting on the source material here. We're not taking political sides on this lawsuit at all but we just want to convey what's in these legal briefs.

Exactly, but the sheer weight of that bipartisan multi-agency legal assault meant the 2024 law was effectively dead on arrival. Okay, so I guess to use a real estate analogy, it's like tearing down a house that failed its foundation inspection and just building a completely new structure from scratch rather than trying to like patch a leaky roof. That's a perfect way to put it.

So Colorado deleted the old software and installed something entirely different with Senate Bill 26299 which takes effect on January 1st, 2027. Right, so what is the new philosophical approach here if they ditch the whole reasonable care thing? So they abandoned that vague prevent all discrimination mandate and they replaced it with a highly rigid notice and

disclosure regime. Instead of forcing you to audit the AI's internal logic beforehand, the new law essentially says you can use the AI but you have to be radically transparent about it.

And you have to give the consumer heavy ammunition to fight back if the AI makes a decision they don't like. Which means we really need to define the battlefield here. Like if I am running a brokerage, I need to know the exact moment I crossed the line into this regulated territory.

So let's look at the new terminology. The state dropped the phrase high-risk AI completely and they replaced it with covered automated decision-making technology or covered ADMT. Right, covered ADMT.

And the definition of that hinges entirely on this concept of a consequential decision. So without getting lost in the weeds of like machine learning architecture, the starting point is this. If you use a system that processes personal data to generate a prediction, a score, or a ranking and that output materially influences a consequential decision, you trigger the law.

And the law explicitly lists the domains where a decision is deemed consequential, right? Like employment, lending, insurance, healthcare. Yes. But the massive red flag for our audience is the explicit inclusion of the lease or purchase of residential real estate in Colorado.

Exactly. That's the trigger. If your AI helps decide who gets a lease or who gets a mortgage, the heavy regulatory machinery just turns on.

I want to challenge that phrase materially influences because honestly, that sounds like a massive legal gray area. It can be, yeah. Let's say I use an AI tool to rapidly sort through like 100 housing applications, and the AI filters out 90 of them based on its own proprietary risk score.

And I, the human broker, I only manually review the top 10. Did the AI materially influence those 90 that just got bumped to the bottom? Yes, absolutely. If the algorithm dictates the primary filtering mechanism that removes a consumer from consideration, that is a material influence.

I mean, the system made a substantive contribution to the outcome, right? It made the cut. Right. But what about just like mundane tech? Say a property manager is using a basic Microsoft Excel spreadsheet to track rent, or maybe they just have a simple chat bot on their website that regurgitates business hours.

Are they suddenly drowning in compliance paperwork? No, no. The law does provide explicit exemptions to prevent that exact scenario. Oh, thank goodness.

Yeah. Technologists that don't utilize machine learning to make complex predictions, they're exempt. Spreadsheets, standard databases, spam filters, anti-malware, firewalls, even general purpose chatbots that cannot make decisions about a person's qualifications, those are entirely carved out.

You really only enter the danger zone when the tech is making an evaluation of a human being. Okay, that makes sense. So as a starting point, if AI helps make a decision on housing-

like scoring a tenant's lease application or filtering buyers, this bill instantly impacts your practice.

Instantly. So okay, you've crossed that threshold. You purchased the fancy predictive tenant screening software.

What is the very first thing an employing broker has to do under this new law? Okay, so the law classifies the business using the software as a deployer. A deployer. And as a deployer, before you ever feed a single application into that system, you must provide a clear and conspicuous public notice.

Wait, what does that look like practically? Just like a pop-up on the application portal? Pretty much. It needs to be reasonably accessible right at the point where the consumer interacts with you. It has to explicitly state that covered ADMT is being used to assist in the decision making process.

And really importantly, it must direct the consumer to a larger policy explaining how they can get more information about the whole system. Honestly, that feels like the easy part. Just updating the website disclosures, right.

But what happens when the algorithm actually rejects someone? Yeah, that is where things get intense. That triggers the most operationally heavy part of this legislation, the 30-day adverse action rule. Okay, explain that.

So if the AI produces an adverse outcome, meaning it denies a housing application, or maybe it calculates that an applicant should be offered materially less favorable lease terms, a clock starts ticking the exact second that decision is made. If you're running a boutique property management firm with a high volume of turnover, you are probably tearing your hair out hearing about a ticking clock. What exactly has to happen within those 30 days? You have to deliver a highly specific plain language notice to the rejected applicant.

You cannot just say, sorry, your application was denied. Because it's not just a standard rejection letter anymore. Exactly.

You must explain the consequential decision, detail the specific role the AI system played in making that decision, identify the data points the AI relied upon, provide the name of the AI developer, and on top of all that, outline instructions on how the consumer can exercise their rights to fight the decision. That is, I mean, just think about the administrative overhead there. Employing brokers need to urgently revisit their operational guidelines.

Honestly, they probably need to create brand new policies from scratch for how their firm will use AI. Oh, 100%. You can't just like quietly adopt a new tech vendor to screen tenants and fly under the radar anymore.

No, those days are over. And to make matters more serious, you have to retain the records of all these decisions and compliance notices for three full years. Four years.

Yep. Every version updated, the algorithm, every notice sent out, it all has to be archived in case the state comes looking. And remember, the old law had some carve outs for federally regulated entities.

Those are gone. Everyone is on the hook. So, okay, you've slapped this highly detailed 30 day notice on an applicant's rejected file.

You've essentially handed them a target. What ammunition does this new law give the consumer to fire back at that rejection? Consumers now have a standalone set of rights. They have the right to access the specific personal data that AI used.

They have the right to demand corrections if the AI pulled factually inaccurate data. Right. But the absolute biggest weapon they have is the right to demand a meaningful human review and a reconsideration of the adverse decision.

Okay, let's pause there. Meaningful human review. The statute tries to define this by saying the reviewer has to be trained, has the authority to override the system, doesn't just blindly default to the AI's output and considers primary evidence.

Right. That's the legal definition. Right.

But if you, the listener, can figure out what a meaningful human review actually looks like in the trenches of a busy real estate office, please let us know. Because, I mean, our guess is that this means a human must manually relook at a denied application to make sure no errors or hallucinations have occurred. Honestly, that is the core of it.

We know large language models and predictive algorithms hallucinate. Oh, all the time. They can invent facts out of thin air or they misinterpret data like confusing a completely resolved utility dispute with a formal eviction.

Yikes. Right. And the psychological hurdle here is something called automation bias.

It's the human tendency to look at a computer screen, see a complex mathematical score, and just assume the machine is smarter than you are. Oh, sure. The computer says no, so it must be no.

Exactly. The law is trying to force brokers to actively fight that bias. So you cannot just be a rubber stamp for the machine.

Yeah. You have to actually look at the primary evidence, like the pay stubs, the references, the credit history, and perform an independent evaluation. Right.

Which brings up a really critical question of liability. Let's say the human broker doesn't catch the error. Or worse, the AI tool turns out to be systematically biased and violates fair housing laws by discriminating against protected classes.

Who actually takes the fall? Oh, I can hear every broker listening right now thinking, I'm just buying a software license. The tech company built the algorithm, they should be the ones liable. Yeah.

Well, the law anticipates that exact offense, and it completely shuts it down. Yes. Brokers cannot contract away their liability for discrimination.

The bill states explicitly that any contractual clauses that try to indemnify a party for its own acts related to violations of the Colorado Anti-Discrimination Act are void as against public policy. Wow. So if the Division of Real Estate, the DRE, comes knocking because your firm's AI tool is unlawfully screening out applicants based on race or family status, you cannot point your finger at the vendor.

Nope. You deploy the tool, you own the outcome. Now, the state will look at relative responsibility, sure, but you absolutely cannot wash your hands of a liability just because you didn't write the code.

Man, that forces us to look up the supply chain then. If the broker is left holding the bag on the front lines, what exactly is the tech developer's responsibility? Like, what do the creators of these AI tools have to deliver to the real estate professionals who are buying them? Good question. So developers are legally required to provide a comprehensive technical documentation to the deployers.

It acts as what the industry calls a model card. Okay. Let's try to avoid the cliche of calling it a car owner's manual because, let's be honest, we all know nobody actually reads the 50-page terms and conditions when they update their iPhone, right? Very true.

If you hand a busy property manager a highly technical, dense, algorithmic document, it's just going to sit in a drawer. So what is actually in this model card? Okay. Think of the algorithm like a commercial kitchen baking a complex cake, which is the final decision.

Okay. I like cake. Right.

So the deployer is legally required to list out every single ingredient used in that cake. So the developer has to give the broker that exact ingredient list, the categories of training data used, the known limitations of the system, the intended uses, and, crucially, the known inappropriate uses. Oh, I see.

So if the consumer finds out the AI used spoiled, milk-like, inaccurate data, maybe a hallucinated eviction, they have the right to demand the broker bake the cake again manually. Exactly. And the broker can only do that if the developer provided the exact recipe instructions on how to monitor the system in the first place.

That is an excellent way to conceptualize it. And, crucially, the developers are not required to hand over their proprietary source code or their trade secrets. Okay.

The tech companies would never go for that anyway. Right. But they must provide enough transparent operational mechanics so that the broker can actually conduct that meaningful human review.

Okay. So what happens when the system just breaks down? A property manager misses the 30-day window. Or maybe they just fail to provide that super complex notice.

We know the consumer can demand a human review, but can they sue the broker directly under this new AI law for, like, millions of dollars? So this was a massive point of negotiation in the legislature, and the final answer is no. Yeah. There is no private right of action under SB 26189.

A consumer cannot just hire a lawyer and sue a broker directly for violating these specific transparency and disclosure rules. Wait, but they can still sue for actual housing discrimination, right? Oh, absolutely. Existing civil rights lawsuits, fair housing violations, and anti-discrimination claims are entirely untouched.

If the AI discriminates, you will be sued under standard civil rights laws. Got it. But for the purely procedural violations of this new AI act-like, failing to post the public notice on your website enforcement is the exclusive domain of the Colorado Attorney General.

Because the AG treats it as a deceptive trade practice. Correct. Now, there is a temporary safety net built into the system right now.

Generally, the Attorney General is required to provide the business with a 60-day notice and a right to cure the violation before launching a full enforcement action. Okay, so if a boutique firm genuinely messes up their application portal disclosure, the AG sends a letter, and the firm has 60 days to rewrite the policy and just get back into compliance. Assuming a cure is deemed possible by the AG, yes.

Okay. However, that 60-day right to cure is essentially just a set of training wheels, and they come off very quickly. What do you mean? That provision completely expires on January 1st, 2030.

Oh, wow. Yeah. After that date, the Attorney General can just drop the hammer immediately without any warning.

So the state is basically giving the industry a three-year runway to figure out the internal compliance mechanics. But you really have to start now. I mean, the AG is required to adopt specific rules clarifying these post-adverse outcome disclosures by January 4th, 2027.

So over the next several months, we're going to see a flurry of regulatory rulemaking dictating exactly what these 30-day notices need to look like. It is going to require constant vigilance from employing brokers to ensure their firm's policies actually match whatever the AG eventually publishes. Well, synthesizing everything from our discussion today, the rules of the game for using AI in Colorado real estate have been completely rewritten from the ground up.

Entirely. The vague burdens of the 2024 law are dead, they're gone, and they've been replaced by the rigid transparency requirements of SB 26189. If your software uses personal data to output a score that helps you decide who gets a house, you are stepping into a heavily regulated minefield.

You really are. You are committing your firm to strict public disclosures, a mandatory 30-day explanation window for any application rejection, and you are taking on the very heavy burden of ensuring a qualified human is always ready to manually override the machine when it inevitably makes a mistake. And January 1st, 2027 is the absolute deadline.

You cannot wait until then to start auditing the software tools your firm currently relies on. No, definitely not. But before we wrap up, let's just look at the broader horizon for a second.

What is the ultimate implication of all this regulation? You know, it raises a fascinating, almost paradoxical question. How so? Well, real estate firms are adopting AI to achieve massive speed and efficiency, right? Right. But if a broker is legally forced to implement a rigorous, meaningful human review, manually double-checking the primary evidence of denied applications to hunt for algorithmic hallucinations, does the regulatory red tape actually cancel out the efficiency the AI promised in the first place? I mean, if you have to do the work anyway, are you really saving time? That is a great point.

Or, I mean, to look at it from a competitive angle, does this intense regulatory environment create a new gold standard in the industry? What do you mean? Well, instead of viewing compliance as a burden, what if human-verified AI becomes a premium marketing tool? Oh, interesting. Right. A property manager could build immense trust with applicants by advertising that a real human being fiercely audits every single algorithmic decision.

Turning mandatory transparency into a competitive advantage. I like that. It's a very real possibility for the firms that get out ahead of this.

It's all about how you adapt to the new lines being drawn. Well, thank you for joining us for the topic today. Take a hard look at the application portals, the screening tools, and the predictive software your firm is using right now.

Ask your vendors exactly how they plan to provide you with the required data and get your firm's internal policies set long before January 2027. We will see you next time.
Colorado Broker Podcast by Apex Real Estate School